

Faster Reinforcement Learning Using Active Simulators

Vikas Jain
Indian Institute of Technology Kanpur
vksjn18@gmail.com

Theja Tulabandhula
University of Illinois Chicago
tt@theja.org

March 24, 2017

Abstract

In this work, we propose several online methods to build a *learning curriculum* from a given set of target-task-specific training tasks in order to speed up reinforcement learning (RL). These methods can decrease the total training time needed by an RL agent compared to training on the target task from scratch. Unlike traditional transfer learning, we consider creating a sequence from several training tasks in order to provide the most benefit in terms of reducing the total time to train.

Our methods utilize the learning trajectory of the agent on the curriculum tasks seen so far to decide which tasks to train on next. An attractive feature of our methods is that they are weakly coupled to the choice of the RL algorithm as well as the transfer learning method. Further, when there is domain information available, our methods can incorporate such knowledge to further speed up the learning. We experimentally show that these methods can be used to obtain suitable learning curricula that speed up the overall training time on two different domains.

1 Introduction

In transfer reinforcement learning (RL), the knowledge obtained from training on a source task can be leveraged to learn a target task more efficiently [17]. A *source task* is typically similar to and sometimes lesser complex than the target task. When such a source task is not readily available, one typically has to select or build a few candidate tasks from scratch and hope that the total time spent to train on (a subset of) these *training tasks* and then the target is less than the time needed to learn a policy from the target task directly.

Narvekar et al. [10] present preliminary attempts in direction. They define the problem of *curriculum learning* for RL as follows: Design a sequence of tasks (i.e., a curriculum) on which a learning agent learns sequentially by transferring knowledge across the stages of the curriculum, ultimately leading to reduced learning time on the target task.

There are two parts to any solution for curriculum learning: the first is to create a set of source tasks, and the second is to determine a curriculum from these. Though in [10], the authors showed that a learning curriculum can speed up the learning on the target task, they stop short of addressing the process of finding a curriculum from the source tasks. Instead, they focus on creating agent-specific training tasks from the target task. In this work, we address this latter key part of the curriculum learning problem, viz., given a set of training tasks, how can one order them in an online fashion to form a learning curriculum that speeds up the overall training for a target task.

We propose methods for online learning of a curriculum in two scenarios: first, when there is no domain knowledge about the tasks at hand, and second, where there is additional knowledge about

the tasks. In the latter case, such domain information can be encoded in various ways (for instance, in [10], the authors are able to design a parametric model for tasks using domain knowledge). We use the available domain information as features to extrapolate learnability across training tasks, which consequently is used to further speed up curriculum learning. In both settings, our proposed methods are independent of the choice of RL algorithm employed by the agent as well as the choice of the transfer technique used (the latter can depend on the representation maintained by the RL agent, viz., Q -values, policies or the model estimates). We only assume that the RL algorithm and the transfer technique used are compatible with each other. Thus our methods can work with existing RL algorithms and transfer methods introduced in the literature [4, 7].

2 Related Work

Offline curriculum learning in the context of supervised learning has been explored before. In [3], the authors showed that learning using a curriculum has an effect on the rate of convergence of the prediction model. In [12], the authors proposed an approach that processes multiple supervised learning tasks in a sequence and finds the best curriculum of tasks to be learned. This is analogous to our reinforcement learning setting, although the emphasis in [12] is on empirical risk minimization rather than training time. Active learning methods have also been used to find a curriculum of supervised learning tasks in the *lifelong machine learning* setting [13, 11]. However, as mentioned before, all these methods apply to the supervised learning setting and typically need to exploit domain knowledge. In contrast, for reinforcement learning, as we show in our experiments, domain knowledge is not necessary for online curriculum learning and can still lead to faster training of an agent on the target RL task.

Analogous to our work is that of [9], wherein the authors define a student-teacher setting for supervised learning. There, the teacher’s task is to provide a sequence of training data to the supervised learning algorithm such that the latter can quickly learn the correct hypothesis. As we show later, our methods follow the template of presenting a task to the RL agent for a number of steps and then changing the task (for instance, picking a slightly more difficult but transfer-learnable task) such that the agent is steered towards the target task.

Transfer in reinforcement learning addresses the problem of designing schemes for efficient knowledge transfer from a given source task to a given target task [2, 18, 15]. However, the focus here is typically on efficient transfer (which in itself is a hard problem) and not on the selection of source tasks. Certain recent works in multi-task reinforcement learning look at the problem of task selection where the agent is presented with a set of tasks to learn from [20]. Here, the goal is to make the agent learn policies that are simultaneously good for all the tasks. As such, this is very different from our setting where we want our agent to learn a target task successfully in a time efficient manner.

In [14], the authors presented an offline method to find a curriculum by computing transfer measure for each pair of training tasks to form an inter-task transferability model. They then use this model to find a curriculum in a recursive fashion offline. Not only this is computationally intensive, such an approach is bound to fail in our online setting where we measure the success of learning a curriculum based on the total time to train (i.e., the time to learn the curriculum and the time to train on the target task). The computation of the inter-task transferability model requires a significant amount of time and domain knowledge. In that time, one could easily train an RL agent on the target task directly from scratch. Nonetheless, when there is domain information available in the form of features, we are able to build on this approach by proposing an active learning method. This method of ours builds a model similar to that in [14] while benefiting from reduced

computation time. One of the directions for future exploration would be to decide in an online manner, when domain knowledge is helpful and when it is not. We currently do not address this choice. It is intuitive that when the number of training tasks is large compared to the complexity of the target task, it may be helpful to use the domain aware online curriculum learning methods compare to the domain-agnostic ones. To the best of our knowledge, we are the first to present general methods for creating a curriculum of training tasks online in both the domain agnostic and domain aware scenarios.

Finally, note that in theory, designing a curriculum while designing the training tasks in a coupled manner can further improve the total time needed for an agent to learn a target task. Unfortunately, the current state of the art is not yet amenable to a coupled approach. Our work complements the work in [10], which present methods to create a set of training tasks which are related to a target task. A future extension could look into the more complex coupled task design and sequencing. In this work, we reasonably assume a set of training tasks similar to the target task are already present and we only need to create a curriculum and learn on the target task faster than learning on the target directly from scratch.

3 Preliminaries

We introduce a couple of key ingredients that will be useful in defining our problem and solutions.

Transfer in RL: In transfer reinforcement learning, instead of learning on a task \mathcal{M} directly, an agent first learns on a suitable source task \mathcal{M}_{sub} . The learned knowledge (for instance, Q -values) from \mathcal{M}_{sub} is then reused to warm-start learning on the target task \mathcal{M} making learning the latter task faster. Different works propose different ways of transferring the knowledge from the source task \mathcal{M}_{sub} to \mathcal{M} including samples [8], the learned policy [5], model of the environment [4] and the value function [16]. In our work, we demonstrate the performance of our online curriculum learning methods using Q -function transfer. This entails initializing Q -function for the agent working with \mathcal{M} using those obtained for \mathcal{M}_{sub} . In the literature, there are ways to facilitate such a transfer even if the state and action spaces of the two tasks are distinct, as long as there is a cross-task mapping between them [19].

Active learning: We discuss active learning in the well known setting of ordinary least squares (with additive Gaussian noise assumption). In this problem, the objective is to find a model that minimizes the Mean Squared Error (MSE) between the true values and the predicted values. That is, we minimize over θ , the objective $\sum_{t=1}^n (y_t - \theta^T (x_t^T))_1^2$, where $(x_t, y_t)_1^n$ is the training data. An analytical solution to this problem (under certain conditions) is given by $\hat{\theta} = (X^T X)^{-1} X^T y$, where X is a matrix whose rows are indexed by training examples $(x_t^T)_1^n$ ($1 \leq t \leq n$) and y is a column vector indexed by training example ground truth values (y_t) . Further, let θ^* be the true parameter that linearly relates y_t and x_t . Then, the expected Euclidean distance between the true and the estimated parameters is given by $E(\|\hat{\theta} - \theta^*\|^2) = \sigma^2 \text{Tr}[(X^T X)^{-1}]$ where σ^2 is the noise variance and $\text{Tr}[\cdot]$ is the trace operator.

In the active learning variant of this problem [6], we sequentially select training examples to minimize the estimation error. Let's assume that we already have some training points (so X is well defined) and let $A = (X^T X)^{-1}$. If we now choose a training example $(x^T)_1$ and append it to X , then A matrix gets updated as: $\begin{pmatrix} X \\ x^T_1 \end{pmatrix}^T \begin{pmatrix} X \\ x^T_1 \end{pmatrix} = (X^T X) + \begin{pmatrix} x \\ 1 \end{pmatrix} \begin{pmatrix} x \\ 1 \end{pmatrix}^T = A + vv^T$ where $v = \begin{pmatrix} x^T_1 \\ 1 \end{pmatrix}^T$. Since we want to choose the training example such that the estimation error is minimized, we can

write estimation error in terms of the new example:

$$\text{Tr}[(A + vv^T)^{-1}] = \text{Tr}[A] - \frac{v^T A A v}{1 + v^T A v}.$$

And observe that we only need to choose a v that maximizes $\frac{v^T A A v}{1 + v^T A v}$. If we assume that our data is normalized (and hence $\|v\| = 1$), then the maximizing v , \hat{v} , turn out to be the normalized eigenvector of A with the largest eigenvalue. We can thus choose the next training point that is most similar to \hat{v} .

4 Problem Setup

Let the MDP corresponding to the target task be \mathcal{M}^{target} . Let the set of training tasks be $\mathcal{T}_{sub} = \{\mathcal{M}_k^{sub} : 1 \leq k \leq K\}$, where each \mathcal{M}_k^{sub} is an MDP similar to \mathcal{M}^{target} with possibly different state and action spaces. Let the total number of steps (each corresponding to a state transition) allowed for training the RL agent be T .

The objective of curriculum learning is to design an algorithm that, when given an RL agent as an input, constructs a (sub)sequence of tasks from \mathcal{T}_{sub} and trains the agent on them first and then on the \mathcal{M}^{target} to achieve a desired level of performance. And in particular, the total number of steps it takes for designing the curriculum, training the agent on the curriculum and then on the target \mathcal{M}^{target} is lesser than the steps needed to train the agent directly on \mathcal{M}^{target} for the same level of performance. In other words, the algorithm (active simulator) actively chooses and switches the tasks which the RL agent is training on. It is assumed that the RL agent is able to transfer knowledge from previous tasks while starting a new task. The algorithm itself can be agnostic to the exact transfer technique. Let the (sub)sequence of tasks designed by the algorithm be $\mathcal{M}_1^* \rightarrow \mathcal{M}_2^* \rightarrow \dots \rightarrow \mathcal{M}_i^* \rightarrow \mathcal{M}^{target}$ where $\mathcal{M}_j^* \in \mathcal{T}_{sub}$ $1 \leq j \leq i$. Further let the ordered set of curriculum tasks be denoted by $\mathcal{I}_m = (\mathcal{M}_1^*, \mathcal{M}_2^*, \dots, \mathcal{M}_i^*)$.

We hypothesize that learning the target task after transfer from a learned curriculum is faster than learning on the target task alone or by using single stage transfer methods. Intuitively, the algorithms that address online curriculum learning have to find a learning curriculum (a sequence) from these training tasks and make the agent learn on the sequence by transferring knowledge from one task to another, with the ultimate aim of learning the target task in the lowest number of simulation steps. Because these algorithms change the task (*i.e.* the simulation environment) seen by the agent over time, we call them *active simulators*.

The algorithms are given a budget of total T steps within which they need to train the RL agent on the target task, where each step signifies one state–action–reward–state cycle. All our algorithms operate in two phases: curriculum selection and training on training tasks, and training on the target task. The first phase can have a budget of $g(T) < T$ steps that can optionally be provided as an input. The algorithms also have access to an oracle $TLearn(\mathcal{L}, \mathcal{M}, \tau)$, where \mathcal{L} is an RL agent, \mathcal{M} is an RL task and τ is the upper bound on number of steps available for learning. This oracle is an abstraction to the process where the RL agent trains on a task for a given number of steps. At the end of training, the oracle returns the reward accumulated by \mathcal{L} on the task \mathcal{M} that it trained on.

Further, the algorithms have no control over the RL algorithm (say a tabular method or a function approximation based method) that the agent \mathcal{L} employs and treat it as a black box. One of the consequences of this is that we cannot reset the learned knowledge of the agent. To get around this, we assume that our algorithms can clone the agent at any given step to freeze the learned knowledge (more details below). We also assume that the agent \mathcal{L} can first transfer

previously learned knowledge (say in the form of a policy or Q -values or the dynamics) using a transfer method and then learn a given task \mathcal{M} .

5 Active Simulators for Online Curriculum Learning

We now present several active simulator variants under two scenarios, viz., the domain agnostic and the domain aware settings. In the domain agnostic setting, online curriculum building is based on the task specific rewards accumulated and their transformations. In the domain aware setting, curricula are obtained using features that describe each of the tasks. The active simulator methods are: (a) Binary Curriculum Selector, Reward Maximizing Greedy Selector, Local Transfer Maximizing Selector, (d) Active Reward Maximizing Greedy Selector, and (e) Active Local Transfer Maximizing Selector. Among them, the first three are domain agnostic and the latter two are domain aware methods.

5.0.1 Binary Curriculum Selector

In this method (see Algorithm 1), the agent is assigned tasks on which the rewards it has accumulated in the past is less. The intuition is that tasks in which the agent has relatively collected more rewards, are less likely to help with good transfer of knowledge in the future compared to tasks where it has not. In essence, the rewards collected are a proxy for the difficulty levels of the tasks and an implicit assumption that more complex tasks are closer to the target task is being made. The method works as follows. First, the agent is trained sequentially on each task in \mathcal{T}_{sub} for a fixed number of steps τ and rewards accumulated for each task ($\mathcal{M}_i.r$) are logged. Then, \mathcal{T}_{sub} is modified to include those tasks on which the accumulated reward is lesser than the remaining half. This shrinks the size of \mathcal{T}_{sub} by 2. And the process is repeated till \mathcal{T}_{sub} is empty. After this, the agent learns on the target task \mathcal{M}^{target} for remaining number of steps. Note that transfer of knowledge is assumed every time there is a switch in the task assigned. The number of steps spent in training with the source task can be a proportion of the total time (say $0.6T$) and automatically determines τ . Finally, note that because of eliminating half the tasks after each stage, the total number of times the training tasks are switched is $K \log K$.

Algorithm 1: Binary Curriculum Selector

Input: $\mathcal{L}, \tau, \mathcal{T}_{sub} = \{\mathcal{M}_k^{sub}\}_{k=1}^K, \mathcal{M}^{target}, T$
Output: \mathcal{L}

```

1  $g(T) = 0$ 
2 while  $\mathcal{T}_{sub} \neq \phi$  do
3    $\mathcal{L}, \mathcal{M}_i.r \leftarrow TLearn(\mathcal{L}, \mathcal{M}_i, \tau) \forall \mathcal{M}_i \in \mathcal{T}_{sub}$ 
4    $g(T) \leftarrow g(T) + \tau |\mathcal{T}_{sub}|$ 
5   Sort  $\mathcal{T}_{sub}$  in ascending order of  $\mathcal{M}_i.r$ 
6    $\mathcal{T}_{sub} \leftarrow \mathcal{T}_{sub}[1 : \frac{1}{2}|\mathcal{T}_{sub}|]$ 
7 end
8  $\mathcal{L} \leftarrow TLearn(\mathcal{L}, \mathcal{M}^{target}, T - g(T))$ 

```

5.0.2 Reward Maximizing Greedy Selector

In this method (see Algorithm 2), the next task in the curriculum is selected greedily based on the performance of the agent on each task in the set of remaining tasks to choose from. In order to

achieve this, the method needs access to multiple copies (clones) of the RL agent \mathcal{L} so that it can accurately measure the transfer effect. This is because, when the agent trains on a task \mathcal{M}_u and then transfers to train on another task \mathcal{M}_v (for instance, to measure the transferability between the two tasks), the information state of the agent at the end of training on \mathcal{M}_v is different from the information state at the beginning of training. For a third task \mathcal{M}_w , one cannot now compute its transferability with \mathcal{M}_u unless the state of the agent after training on \mathcal{M}_u is preserved. The use of copies helps the method replicate an agent’s knowledge level for computing inter-task transferability reliably. Note that making a copy is a mild assumption compared to having direct access to the internal state of the learning agent.

Thus there are multiple learners: (a) the original agent \mathcal{L} , and (b) virtual learners $\{\mathcal{L}_c\}$ that copy agent \mathcal{L} ’s state. A copy agent \mathcal{L}_c is used to assess inter-task transferability effect starting from the current state of \mathcal{L} . The method then selects the task where the copy agent gets the maximum reward and adds it to the curriculum. The original agent \mathcal{L} then learns from this new task that was added to the curriculum. This process is repeated till all the subtasks are added sequentially into the curriculum. After training on all tasks in the curriculum, the method trains the agent \mathcal{L} on the target task \mathcal{M}^{target} . One key aspect of this method is that this selects the next task based on the sequence of tasks already in the curriculum at every decision epoch.

In Algorithm 2, τ_1 and τ_2 are the number of steps used for evaluating inter-task transferability and learning a task respectively. Procedure $\text{Eval}(\mathcal{L}_c, \mathcal{M}_i, \tau_1)$ outputs the reward accumulated on task \mathcal{M}_i by \mathcal{L}_c in τ_1 number of steps. Note that, this procedure does not alter the original RL agent \mathcal{L} . The length of the curriculum obtained in this method is equal to the number of training tasks K .

Algorithm 2: Reward Maximizing Greedy Selector

Input: \mathcal{L} , τ_1 , τ_2 , $\mathcal{T}_{sub} = \{\mathcal{M}_k^{sub}\}_{k=1}^K$, \mathcal{M}^{target} , T
Output: \mathcal{L}

```

1  $g(T) \leftarrow 0$ 
2 while  $\mathcal{T}_{sub} \neq \phi$  do
3    $\mathcal{L}_c \leftarrow \text{copy}(\mathcal{L})$ 
4    $\mathcal{M}_i.r \leftarrow \text{Eval}(\mathcal{L}_c, \mathcal{M}_i, \tau_1) \ \forall \mathcal{M}_i \in \mathcal{T}_{sub}$ 
5    $\mathcal{M}^* \leftarrow \arg \max_{\mathcal{M}_i \in \mathcal{T}_{sub}} \mathcal{M}_i.r$ 
6    $\mathcal{L} \leftarrow \text{Tlearn}(\mathcal{L}, \mathcal{M}^*, \tau_2)$ 
7    $g(T) \leftarrow g(T) + \tau_1 |\mathcal{T}_{sub}| + \tau_2$ 
8    $\mathcal{T}_{sub} \leftarrow \mathcal{T}_{sub} \setminus \mathcal{M}^*$ 
9 end
10  $\mathcal{L} \leftarrow \text{Tlearn}(\mathcal{L}, \mathcal{M}^{target}, T - g(T))$ 
```

5.0.3 Local Transfer Maximizing Selector

In this method (see Algorithm 3), the effect of transfer from one task to another task is calculated for each task pair in \mathcal{T}_{sub} as well as for pair of tasks where the first task is in \mathcal{T}_{sub} and the second task is \mathcal{M}_{target} . Based on the $K \times (K + 1)$ numbers computed, the methods designs a curriculum that maximizes local transfer. That is, the curriculum is obtained by starting from \mathcal{M}_{target} and traversing in the reverse order. In particular, a predecessor task with which the transfer effect to \mathcal{M}_{target} is the highest is first selected. Then, the process is repeated recursively and a sequence of tasks $\mathcal{M}_1^* \rightarrow \mathcal{M}_2^* \rightarrow \dots \rightarrow \mathcal{M}_{target}$ is obtained. Then the method assigns this curriculum to the

RL agent \mathcal{L} . As in the previous method, copies of the agent \mathcal{L} are used to accurately capture the transfer effect.

In Algorithm 3, the quantity F is the inter-task transferability matrix where each entry $F[i][j] = \text{TransferMeasure}(\mathcal{L}_c, \mathcal{M}_i, \mathcal{M}_j, \tau_1)$ stores the reward accumulated while training on task \mathcal{M}_j after transferring knowledge (a learned policy or Q -values etc.) from experience with task \mathcal{M}_i and starting with task \mathcal{M}_j . Note that the copy agent \mathcal{L}_c also follows \mathcal{L} 's transfer method while capturing transferred knowledge. The method spends τ_1 steps for computing the transfer effect for each pair of tasks. The length of the curriculum obtained in this method is K , the size of the training task set \mathcal{T}_{sub} . This method is in principle an offline method similar to [14] but without taking domain knowledge in consideration.

Algorithm 3: Local Transfer Maximizing Selector

Input: \mathcal{L} , τ_1 , τ_2 , $\mathcal{T}_{sub} = \{\mathcal{M}_{sub_k}^{sub}\}_{k=1}^K$, \mathcal{M}^{target} , T
Output: \mathcal{L}

```

1  $F[K][K+1] \leftarrow 0$ ,  $g(T) = 0$ 
2 for  $\mathcal{M}_i \in \mathcal{T}_{sub}$  do
3   for  $\mathcal{M}_j \in \mathcal{T}_{sub} \cup \mathcal{T}_{target}$  do
4      $\mathcal{L}_c \leftarrow \text{copy}(\mathcal{L})$ 
5      $F[i][j] \leftarrow \text{TransferMeasure}(\mathcal{L}_c, \mathcal{M}_i, \mathcal{M}_j, \tau_1)$ 
6      $g(T) \leftarrow g(T) + \tau_1$ 
7   end
8 end
9  $C \leftarrow \{\mathcal{M}^{target}\}$ ,  $NextTask \leftarrow \mathcal{M}_{target}$ 
10 while  $\mathcal{T}_{sub} \neq \phi$  do
11    $\mathcal{M}^* \leftarrow \arg \max_{\mathcal{M}_i \in \mathcal{T}_{sub}} F[\mathcal{M}_i][NextTask]$ 
12    $C \leftarrow C.add(\mathcal{M}^*)$ 
13    $\mathcal{T}_{sub} \leftarrow \mathcal{T}_{sub} \setminus \mathcal{M}^*$ 
14    $NextTask \leftarrow \mathcal{M}^*$ 
15 end
16  $C \leftarrow \text{Reverse}(C)$ 
17 for  $\mathcal{M}_i \in C \setminus \{\mathcal{M}^{target}\}$  do
18    $\mathcal{L} \leftarrow \text{TLearn}(\mathcal{L}, \mathcal{M}_i, \tau_2)$ 
19    $g(T) \leftarrow g(T) + \tau_2$ 
20 end
21  $\mathcal{L} \leftarrow \text{TLearn}(\mathcal{L}, \mathcal{M}^{target}, T - g(T))$ 

```

5.0.4 Active Reward Maximizing Greedy Selector

This is the first of the two domain aware methods we propose. In this setting, we are also given feature vectors $f_i = [f_{i1}, f_{i2}, \dots, f_{in}]$ for each task $\mathcal{M}_i \in \mathcal{T}_{sub}$.

We will adapt the Reward Maximizing Greedy Selector method described before (see Algorithm 2) to this information rich setting in the following way. Consider the partial list of tasks (\mathcal{M}^*) at a particular iteration m in the while loop and denote it by \mathcal{I}_m . When we need to decide on the next task in the curriculum at step $m+1$, we will avoid evaluating the reward collected (see line 4 in Algorithm 2) for every possible task under consideration.

In other words, instead of finding the transferability measure for each sequence-task pair (the

sequence being \mathcal{I}_m and the task being one of the remaining tasks in \mathcal{T}_{sub} at step $m + 1$), the method will now use the task specific features to create feature vectors for every pair of tasks and build a transferability prediction model. In particular, the method will apply the active learning technique described in Section 3 to build a regression model and predict the transferability measure for sequence-task pairs. Further, it will relearn the model at each step m .

Let the curriculum sequence at step m be $\mathcal{I}_m = \{\mathcal{M}_1^*, \mathcal{M}_2^*, \dots, \mathcal{M}_{m-1}^*\}$. Feature vectors for the sequence as well as the tasks that are not yet in the curriculum are designed as follows. The feature vector $f^{\mathcal{I}_m}$ for sequence \mathcal{I}_m is computed as: $f_k^{\mathcal{I}_m} = \frac{1}{m-1} \sum_{i=0}^{m-1} f_{ik}^*$, where f_i^* is feature of task \mathcal{M}_i^* . For each remaining task in \mathcal{T}_{sub} (say \mathcal{M}_j), the sequence-task pair feature vector $f^{\mathcal{I}_m j}$ is consequently defined as: $f_k^{\mathcal{I}_m j} = \frac{f_{\mathcal{I}_m k} - f_{jk}}{\max(f_{\mathcal{I}_m k}, \epsilon)}$, where $\epsilon > 0$ is a small positive constant.

To find a new task to be added to the curriculum, a regression model is learned at each step. At step m , the method learns a regression model from \mathcal{I}_m and the remaining tasks in \mathcal{T}_{sub} . Then, it picks that task in \mathcal{T}_{sub} which has maximum value of transferability measure. The length of the curriculum obtained in this method is $\leq n$, because one may not include a particular task based on the regression predictions.

5.0.5 Active Local Transfer Maximizing Selector

The second domain aware method extends Local Transfer Maximizing Selector (Algorithm 3) as follows. Instead of finding transferability measure $F[i][j]$ for every pair of tasks $\forall i, j$, it uses the features of the tasks to form pair wise feature vectors and learn a regression model \hat{F} that can predict the entries of matrix F . Again the active learning technique can be used to build such a model using limited number of task pairs (i, j) , speeding up the curriculum design in the process.

The feature vector f^{ij} of a task pair is computed accounting for the similarity of the two tasks' features. The k -th element f_k^{ij} of the vector f^{ij} is defined as: $f_k^{ij} = \frac{f_{ik} - f_{jk}}{\max(f_{ik}, \epsilon)}$, where $\epsilon > 0$ is a small constant. The regression model \hat{F} is used instead of inter-task transferability matrix F and the algorithm function very similarly to Algorithm 3. The length of the curriculum obtained in this method can also be $\leq n$.

We conclude by summarizing the differences between the two settings. Firstly, since there is no task representation available to domain agnostic methods, ideally the transfer performance needs to be measured for all task pairs (or sequence-task pairs) and based on this information, the next task of the curriculum should be chosen. On the other hand, in the domain aware scenarios, one can utilize the task features to model and predict the task pair (or sequence-task pair) transfer effects and avoid computing them for all pairs. Thus, tools such as active learning can be used (as shown above) to reduce the number of pairs required to learn a transfer effect model, leading to potential gains in the number of steps needed while training on the curriculum of training tasks.

Secondly, in the case of redundant tasks in \mathcal{T}_{sub} , domain independent methods cannot easily prune the tasks. In domain aware scenario, where task features are available, diversity scores $d(\mathcal{M}, \mathcal{I}_m)$ which is the measure of diversity (inverse of similarity) between task \mathcal{M} and the tasks in task-sequence \mathcal{I}_m can be theoretically computed to remove the redundancy.

6 Experiments

We verified the effectiveness of our methods on two domains: (a) a maze environment [1], and (b) a grid world environment (see Figure 1). We found that the curricula selected by our approaches indeed outperform the alternative (training on the target) in all cases. In the domain aware setting, we also demonstrate that active learning methods can speed up training on the target task.

Both domains require an agent (employing the Q -learning algorithm) to find an optimal policy for reaching a goal state. The reward value for reaching the goal state is +1. We set the discount factor (γ) to be 0.9 and learning rate (α) to be 0.6. In the maze task, the start position of the agent is sampled uniformly at random from the set of feasible states while in the grid world, the start position is fixed. The episode ends when the agent reaches the goal position.




Figure 1: *Left*: target task in the maze environment with goal position (*green*), feasible positions (*gray*) and blocked positions (*black*). *Right*: target task in the grid world with goal position (*red*), feasible positions (*white*), blocked positions (*black*) and the start position (*yellow*).

Training Tasks: We construct training tasks for the above two target tasks using methods similar to [10] (see Figure 2). In the maze task, since the start position of the agent is not fixed, the source tasks are created by reducing the state space of the target task while keeping the goal position same. In the grid world, since the start position of the agent is also given, the source tasks are created by changing the start position of the agent while keeping the rest of the task fixed.




Figure 2: Training tasks for the maze (*top*) and grid world (*bottom*) tasks.

We now present both quantitative and qualitative results of performance. In particular, we report the convergence time of our methods as well as the reward accumulated. And we also comment on the curricula selected and the final policy learned.

Quantitative Results: As a *baseline*, we consider the agent learning the target task from scratch. We plot the total reward accumulated by our methods as well as the baseline until the Q -function converges on the target task (with Monte Carlo averaging over 30 runs). Figure 3(*Left*) shows the number of steps required for the convergence as well as the reward accumulated on both the domains.

Qualitative Results: For both environments, we consider 4 source tasks leading to 24 possible curricula of full length. We calculate the number of steps required by each curriculum for convergence on the target task. Figure 4 shows the convergence time of each curriculum as compared to the baseline and the curriculum selected by our methods for both the domains. The final policies learned for both the domains using our methods are shown in the Figure 3(*Right*).

Domain Aware: It can be seen in Figure 3 that the Local Transfer Maximizing Selector [14] is taking larger number of steps for training. We utilize the active learning setting described in section 3 to compute the inter-task transferability matrix using lesser number of steps. In the grid




Figure 3: *Left*: Reward accumulated over the learning task with the number of steps required for convergence for the each proposed method (abbreviated in the legend) and the baseline on the maze (*top-left*) and grid world (*bottom-left*) tasks. *Right*: State-Action table learned after convergence for the maze (*top-right*) and the grid world (*bottom-right*) tasks. Darker shades represent numerically higher values.




Figure 4: Number of steps required for convergence per curriculum (indexed from 0 to 23), compared with the baseline (red) for the maze (*left*) and the grid world (*right*) tasks. The curricula represent all permutations of tasks in Figure 2, ordered lexicographically. Yellow squares represent the most frequent (mode) curricula selected by our methods.

world, the features f_i is defined as the distance between the goal and the start position and f^{ij} is defined as $f^i - f^j$. In the maze world, we directly define inter-task features f^{ij} to be the number of overlapping states between the two tasks divided by the number of states in the task \mathcal{M}_i . Figure 5 shows that the number of steps reduce significantly while the performance and the curriculum selected remain the same.

7 Conclusion

In this paper, we have defined the online curriculum learning problem as a way to speed up the learning of reinforcement learning agents on a target task. We designed several algorithms for two broad settings. In the first, no task information is available and in the second, task dependent features are available. Our algorithms were able to choose curricula and train reinforcement learning agents more efficiently compared to training the agents directly on the target tasks for two different




Figure 5: Comparison of Active LTMS method with LTMS method and baseline in Figure 3 on the maze (left) and grid world (right) tasks.

domains. Statistical and computational guarantees for the proposed algorithms are left for future work.

References

- [1] M. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda. Vision-based behavior acquisition for a shooting robot by using a reinforcement learning. In *Proc. of IAPR/IEEE Workshop on Visual Behaviors*, pages 112–118. Citeseer, 1994.
- [2] A. Barreto, R. Munos, T. Schaul, and D. Silver. Successor features for transfer in reinforcement learning. *arXiv preprint arXiv:1606.05312*, 2016.
- [3] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pages 41–48. ACM, 2009.
- [4] A. Fachantidis, I. Partalas, G. Tsoumakas, and I. Vlahavas. Transferring task models in reinforcement learning agents. *Neurocomputing*, 107:23–32, 2013.
- [5] F. Fernández and M. Veloso. Probabilistic policy reuse in a reinforcement learning agent. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 720–727. ACM, 2006.
- [6] T. Jaakkola. course materials for 6.867 machine learning, fall 2006. mit opencourseware. (<http://ocw.mit.edu/>) Massachusetts Institute of Technology. Downloaded on [10 June 2016].
- [7] A. Lazaric. Transfer in reinforcement learning: a framework and a survey. In *Reinforcement Learning*, pages 143–173. Springer, 2012.
- [8] A. Lazaric, M. Restelli, and A. Bonarini. Transfer of samples in batch reinforcement learning. In *Proceedings of the 25th international conference on Machine learning*, pages 544–551. ACM, 2008.
- [9] J. Liu and X. Zhu. The teaching dimension of linear learners. *Journal of Machine Learning Research*, 17(162):1–25, 2016.
- [10] S. Narvekar, J. Sinapov, M. Leonetti, and P. Stone. Source task creation for curriculum learning. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 566–574. International Foundation for Autonomous Agents and Multiagent Systems, 2016.

- [11] P. Pang, J. An, J. Zhao, X. Li, T. Ban, D. Inoue, and H. Sarrafzadeh. Smart task orderings for active online multitask learning. In *Proceedings of SIAM International Conference on Data Mining*, 2014.
- [12] A. Pentina, V. Sharmanska, and C. H. Lampert. Curriculum learning of multiple tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5492–5500, 2015.
- [13] P. Ruvolo and E. Eaton. Active task selection for lifelong machine learning. In *AAAI*, 2013.
- [14] J. Sinapov, S. Narvekar, M. Leonetti, and P. Stone. Learning inter-task transferability in the absence of target task samples. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 725–733. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [15] M. E. Taylor, G. Kuhlmann, and P. Stone. Autonomous transfer for reinforcement learning. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, pages 283–290. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [16] M. E. Taylor and P. Stone. Behavior transfer for value-function-based reinforcement learning. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 53–59. ACM, 2005.
- [17] M. E. Taylor and P. Stone. Transfer learning for reinforcement learning domains: A survey. *Journal of Machine Learning Research*, 10(Jul):1633–1685, 2009.
- [18] M. E. Taylor and P. Stone. An introduction to intertask transfer for reinforcement learning. *AI Magazine*, 32(1):15, 2011.
- [19] M. E. Taylor, S. Whiteson, and P. Stone. Transfer via inter-task mappings in policy search reinforcement learning. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, page 37. ACM, 2007.
- [20] A. Wilson, A. Fern, S. Ray, and P. Tadepalli. Multi-task reinforcement learning: a hierarchical bayesian approach. In *Proceedings of the 24th international conference on Machine learning*, pages 1015–1022. ACM, 2007.